


PCTWORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶: G06F 11/00, 1/00	A1	(11) International Publication Number: WO 99/38076 (43) International Publication Date: 29 July 1999 (29.07.99)
(21) International Application Number: PCT/US99/01444 (22) International Filing Date: 21 January 1999 (21.01.99) (30) Priority Data: 09/010,939 22 January 1998 (22.01.98) US (71) Applicant: SYMANTEC CORPORATION [US/US]; 10201 Torre Avenue, Cupertino, CA 95014 (US). (72) Inventor: MILLARD, John; 915 N Avenue 64, Los Angeles, CA 90042 (US). (74) Agents: RADLO, Edward, J. et al.; Fenwick & West LLP, Two Palo Alto Square, Palo Alto, CA 94306 (US).		(81) Designated States: CA, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>

(54) Title: COMPUTER FILE INTEGRITY VERIFICATION**(57) Abstract**

System and method for verifying the integrity of contents within a computer file (5). A security value S is stored within the file (5). A verification function f is applied against the entire contents of the file (5) including S, where f is a function of S. Results R of the applying step are compared against a preselected value r, where r is not stored within the file (5). When R equals r, a determination is made that the file (5) has not been modified. f is typically a distributive invertible function such as the Cyclic Redundancy Check (CRC) function known as modulo p, where p is a prime number and is one bit greater than the length of S. Typically, the value of r is zero. Before executing the verification function f, a check generating program is first executed. This check generating program is executed by a computer (4) that is remote from the file (5), further enhancing the security of the system.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

Description

COMPUTER FILE INTEGRITY VERIFICATION

Related Application

A related patent application is U.S. patent application serial number 08/893,934 filed July 15, 1997 and entitled "Secure Access to Software Modules", which patent application has the same inventor and assignee as the present application and is hereby incorporated by reference in its entirety into the present patent application.

Technical Field

This invention pertains to the field of verifying the integrity of the contents of computer files.

5 Background Art

Figure 1 illustrates a generic prior art system for verifying the integrity of computer data 1. File 5 is associated with a digital computer 2, and contains data 1 and a security value S stored in location 7 within file 5. File 5 is
10 accessible to a central processing unit (CPU) 3 of the computer 2. CPU 3 executes program instructions on behalf of the computer 2.

In a first embodiment of the prior art, CPU 3 applies a checksum function against data 1 within file 5, in which all the
15 bytes of file 5, except for the security value S, are added together. This sum is then compared with the security value S. If these two values match, data 1 is deemed not to have been modified, maliciously or otherwise. In this checksum embodiment, the security value S could be stored in a location
20 that is not part of file 5 but is accessible thereto. The problem with this method of file verification is that it is easy to rechecksum file 5 if file 5 has been changed for malicious purposes. For example, a hacker could rather easily maliciously change the contents of data 1, then recompute security value S
25 to correspond with the changed data 1, thus lulling the user of

the computer 2 into thinking that the contents of data 1 had not actually changed. This method of data verification can be easily and regrettably reversed engineered to show that the function being used by CPU 3 is a checksum function.

5 In a second embodiment of the prior art, CPU 3 uses a cyclic redundancy check (CRC) function against the data 1 within file 5. In this embodiment, a CRC of all the bytes in the file 5, not including security value S, is computed by CPU 3 and stored as security value S. Again, S is easy to recompute if
10 the hacker knows that a CRC is being used, and this method of data verification can be relatively easily reverse engineered to show that a CRC is in fact being used.

In a third embodiment of the prior art, CPU 3 uses a cryptographically secure hash function, such as MD5, to create a
15 message digest of the data 1 within file 5, and stores the resulting output hash value in file 5 as security value S. MD5 is described in Schneier, Applied Cryptography, Second Edition (John Wiley & Sons 1996), pp. 436-441, U.S.A. As with the first two prior art embodiments described above, security value S is
20 not used by the hash function when the function is executed. As before, it is easy to recompute the security value S if the hacker knows that MD5 or another hash function is being used, and then to surreptitiously replace the security value S within file 5; and it is easy for the hacker to determine which
25 function is being used.

McNamara, John E., Technical Aspects of Data Communication, 2nd Ed. 1982, Digital Equipment Corporation, U.S.A., pp. 110-122, describes a Cyclic Redundancy Check (CRC) function that is useful in the present invention.

30 Ore, Oystein, Number Theory and Its History, 1976, Gudrun Ore, U.S.A., pp. 124-129, discuss mathematical problems having two unknowns.

Disclosure of Invention

The present invention is a system and method for verifying
35 the integrity of contents within a computer file (5). The

method comprises the steps of storing a security value S within the file (5), applying a verification function f against the entire contents of the file (5) including S, where f is a function of S; comparing results R of the applying step against
5 a preselected value r, where r is not stored within the file (5); and, when R equals r, determining that the file (5) has not been modified.

Brief Description of the Drawings

These and other more detailed and specific objects and
10 features of the present invention are more fully disclosed in the following specification, reference being had to the accompanying drawings, in which:

Fig. 1 is a block diagram illustrating prior art.

Fig. 2 is a block diagram illustrating the present
15 invention.

Fig. 3 is a flow chart illustrating the operation of the present invention.

Detailed Description of the Preferred Embodiments

In this application, reference numeral 7 pertains to the
20 location of the security value, while the letter S pertains to the numerical value of the security value. Similarly, reference numeral 13 pertains to the location of the residual value, while the letter r pertains to the numerical value of the residual value. As used in this application, "data" can be anything that
25 is storable within a digital computer. Thus, data can include computer programs, information that the user wishes to store or manipulate, etc.

The operation of the present invention can be understood by examining Figure 3. At step 31, the file integrity verification
30 function f is selected by the users of computers 2 and/or 4. In the present invention, the file verification function f can be any distributive invertible function. As used in this specification and claims, distributiveness pertains to additive distributiveness. Thus, a function f that is additive
35 distributive is one that satisfies the equation:

$$f(N)+f(S)=f(N+S)$$

where N and S are any variables.

As used in this specification and claims, an invertible function $f(N,S)=r$ is a function where, given r and given N, one
5 can calculate S. r is the residual value that results when f is applied to N and S.

An example of a distributive invertible function that is useful in the present invention is the Cyclic Redundancy Check (CRC) function that is described in McNamara, supra. This CRC
10 function is the modulo p function, where p is somewhat loosely referred to as a "polynomial". The division by p is performed over a Galois field to expedite the calculations. The CRC function is an excellent choice for use in the present invention, because it is not obvious that it is an invertible
15 function; and, even if one does realize that it is an invertible function, it is not obvious how to invert it. These features enhance security, because they make it harder for hackers to determine the type of security being used.

For the CRC function generally, $(N+S2^n) \bmod(p)=R$. In this
20 case, the residual value R is the remainder that is left over after the division by p. For the purposes of this invention, the quotient is irrelevant and is discarded.

For a CRC function generally, p is not normally a prime number, because it is desired to obtain a gratuitous parity
25 check along with the calculation of modulo p. In the present invention, on the other hand, we do want p to be a prime number. p is selected to be one bit longer than the size of the security value S. The first and last bits of p must be a 1. The first bit of p must be a 1 so as to preserve the desired length of p.
30 The last bit of p must be a 1 to guarantee that p really is a prime number. p can be any random prime number satisfying these criteria.

At step 32, the target (desired) residual value r obtained by applying f against the contents of file 5 is selected, and is
35 stored at a location 13 that is accessible to CPU 3 but is not part of file 5. The value of r is typically selected to be zero. The size of location 13 is typically 64 bits.

In step 33, a check generating program calculates the value of security value S , where f is a function of S , based upon the entire contents of file 5. For security purposes, this calculation is performed by a computer 4 which is remote from user computer 2. In step 34, computer 4 stores S at a preselected location 7 within file 5. Location 7 should be embedded within file 5, and not at its beginning or end. At step 35, CPU 3 executes verification function f , applying it against the entire contents of file 5, including all data 1, the computer code 9 representing verification function f if said code is present within file 5, and the security value S itself (see Fig. 2). It is preferable for the preferred embodiment where f is a CRC function (and perhaps for other functions f) for CPU 3 to begin the processing at some start address 10 that is neither at the beginning nor the end of file 5. The processing thus proceeds in a downward direction from the dashed line in Fig. 2, wraps around the end of file 5, and ends just before where it started at start address 10. If the processing were to start at the beginning of file 5, it would be easier for a hacker to defeat the security. Start address 10 should not coincide with the beginning address or the ending address of location 7.

At step 36, the results R obtained by CPU 3 applying function f against file 5 are compared against the target residual value r . When these two numbers are equal, CPU 3 determines that the contents of file 5 have not been modified. This information can be conveyed in the form of a signal 37 that is sent to the user via user interface 15, which may be, for example, a monitor. If, on the other hand, R does not equal r , then CPU 3 determines that the contents of file 5 have been modified, maliciously or otherwise. This information can be signaled to the user by signal 38 over user interface 15. This determination step, like all of the other steps of the present invention, can be implemented in hardware, firmware, software, or any combination thereof.

If a hacker were to modify any of the contents of file 5, including data 1, security value S , and/or verification code 9,

applying the verification function f against file 5 would result in R not equaling r . The user would be flagged by signal 38 that a modification had occurred and thus the contents of file 5 had been compromised.

5 The operation of the check generating function will now be described. This function is executed by computer 4. Computer 4 can be controlled, for example, by the vendor of a software package that is shipped and used by a number of user computers 2. The purpose of the check generating function is to calculate
10 the security value S that is stored within file 5 of each user computer 2. The user is allowed to change the contents of his file 5 after security value S has been stored therewithin, but if he does so and wishes to maintain security, it is necessary for computer 4 to re-execute the check generating function, thus
15 recalculating security value S , and then to re-store S within file 5.

 The check generating function is the same function chosen to be verification function f . Since the preferred function for verification function f is the CRC function, a CRC function will
20 be described herein. The result R of applying a CRC function is a remainder when dividing a sequence of bytes by another sequence of bytes, performing said division while treating the bytes as polynomials over a Galois field based on 2^m , where m is the size of the CRC residual value R . The CRC value R depends
25 sensitively on the entire contents of file 5, but it is a linear function and can be inverted, although not obviously. There are optimized methods of calculating CRC's that do not require actually doing long division in polynomial fields, thus making the implementation quicker. These optimized methods, which are
30 described in the reference cited in endnote 7 of McNamara, supra, make it hard to determine that it is even a CRC that is taking place, thus making the lives of hackers more difficult.

 The check generating program generates security value S as follows. The goal is to calculate that security value S that
35 will cause the actual residual value R of applying verification function f against file 5 to equal the target residual value r . In the following mathematical description, N is the aggregated

entire contents of file 5 treated as one very large number. This value is normalized, so that the first bytes checked by verification program f (i.e., those bytes commencing at start address 10) are the beginning of said very large number N .

5 Thus, N reflects what the CRC function f is actually executing against, not the native order of file 5.

R is the residual value for file 5. S is the security value inserted at location 7 within file 5. n is the position of the security value S from the end of file 5, so that, treated
10 as a number from the point of view of the CRC calculation, the security value is equal to $S2^n$. p is a number called the CRC polynomial that is used to generate the CRC residual value R . r is the target CRC residual value for file 5. As noted above, in the case of the CRC function, all arithmetic is done on
15 polynomials in a Galois field based upon 2^m . For Galois polynomial arithmetic, "+" and "-" are equivalent to each other and are equivalent to the "exclusive or" (XOR) operation.

The initial CRC residual value for file 5:

$R = N \bmod(p)$, by the definition of CRC.

20 The goal is that R equals r , which will normally require modifying S . We achieve this goal by noting that if we modify N by changing S , then we have:

$$N_{\text{new}} = N + S2^n$$

$$R_{\text{new}} = N_{\text{new}} \bmod(p)$$

25 Since mod is a linear function, we have:

$$R_S = S2^n \bmod(p), \text{ where } R_S \text{ is the residue of } S.$$

$$R_{\text{new}} = (N \bmod(p) + S2^n \bmod(p)) \bmod(p) = (R + R_S) \bmod(p)$$

Since "+" is a XOR in this field, and R and R_S are both smaller than p , being already both remainders, this expression
30 reduces to:

$$R_{\text{new}} = R + R_S$$

The desired remainder R_S , is

$$R_S = r - R \text{ because we want } R_{\text{new}} \text{ to be equal to } r.$$

Now we have, by the definition of R_S as given above:

35 $S2^n \bmod(p) = r - R$, and we need to calculate S .

$$R_S \bmod(p) = r - R \text{ implies that:}$$

$pA+r-R=R_S$ for some integer A, so we have:

$$pA+r-R=S2^n$$

Rearranging while remembering that "+" and "-" are equivalent in Galois field 2^m arithmetic:

5 $pA+S2^n=r-R$

This is an equation in two unknowns S and A, requiring a solution for both S and A in integers (Diophantine equation). This can be solved using Euclid's extended algorithm (see Ore, supra). The only requirement for this to have a solution is that
10 p be primitive (the equivalent of a prime number in a Galois field), which ensures that the algorithm does not lead to an indeterminate solution.

We need only S, and are uninterested in A. Computer 4 inserts S into location 7 in the original file 5, preferably by
15 using an XOR operation in case the original contents at location 7 were not zero. We now have our desired result. The CRC residual value R of file 5 is the desired target value r.

The present invention offers the following advantages over the prior art:

- 20 • The file integrity check value (security value S) is not stored within file 5 in a direct way, but file 5 is still self-contained. By this is meant that all of the verification is done using file 5 itself.
- 25 • Reverse engineering the verification function f does not directly help with changing the security value S in the file 5 to subvert the integrity check when the file 5 is maliciously modified. This is because a hacker does not know where the security value S is within file 5, or how to change security value S to get the desired result.
- 30 • There is no computer code within file 5 that does the calculation of the needed security value S, because the calculation of the security value S is performed by computer 4 at a remote location.
- 35 • The verification of file 5 integrity can be done by a relatively simple CRC function.

- A hacker who wishes to maliciously modify data 1 within file 5 wishes for a safe modification of the verification program f. He wants to disable security safeguards of the program f without totally breaking the program f, because he does not want the user to know what is going on. In the present invention, there is no indication within the verification program f where the file 5 can be safely modified to alter the residual value R to insure that R continues to equal r.

The above description is included to illustrate the operation of the preferred embodiments and is not meant to limit the scope of the invention. The scope of the invention is to be limited only by the following claims. From the above discussion, many variations will be apparent to one skilled in the art that would yet be encompassed by the spirit and scope of the present invention.

What is claimed is:

Claims

1. A method for verifying the integrity of contents within a computer file, said method comprising the steps of:

storing a security value S within the file;

5 applying a verification function f against the entire contents of the file including S, where f is a function of S, to obtain results R;

comparing R against a preselected value r, where r is not stored within the file; and

10 when R equals r, determining that the file has not been modified.

2. The method of claim 1, wherein f is a distributive invertible function.

3. The method of claim 2, wherein f is the cyclic
15 redundancy check function known as modulo p, p being a "polynomial".

4. The method of claim 3, wherein p is a prime number.

5. The method of claim 4, wherein the length of p is one bit greater than the length of S.

20 6. The method of claim 1, wherein the value of r is zero.

7. The method of claim 1, wherein a computer remote from the file executes a check generating program; and

the check generating program calculates S and stores S within a preselected location within the file.

25 8. The method of claim 7 where S is calculated to be that number that will force R to be equal to r.

9. The method of claim 7, wherein, subsequent to the final step of claim 7, a user changes the contents of the file; and

the check generating program recalculates and restores
5 S based upon the changed contents of the file, in order to preserve the integrity of the file.

10. Apparatus for safeguarding the integrity of contents within a computer file, said apparatus comprising:

a file associated with a first computer and containing
10 data;

a second computer located remote from the file, said second computer adapted to calculate a security value S, where S is based upon all the contents of the file, and to store S within a preselected location within the file;

15 processing means for applying a verification function f against the entire contents of the file including S, where f is a function of S, to obtain results R; and

means for comparing R against a preselected value r, where r is not stored within the file; and

20 means for determining that the file has not been modified, when R equals r.

11. A computer readable medium for storing a computer program that verifies the integrity of contents within a computer file by following the steps of:


25 applying a verification function f against entire contents of the file including a prestored security value S, where f is a function of S, to obtain results R;

comparing R against a preselected value r, where r is not stored within the file; and

30 when R equals r, determining that the file has not been modified.

1/3

FIG 1
PRIOR ART



2/3

FIG. 2

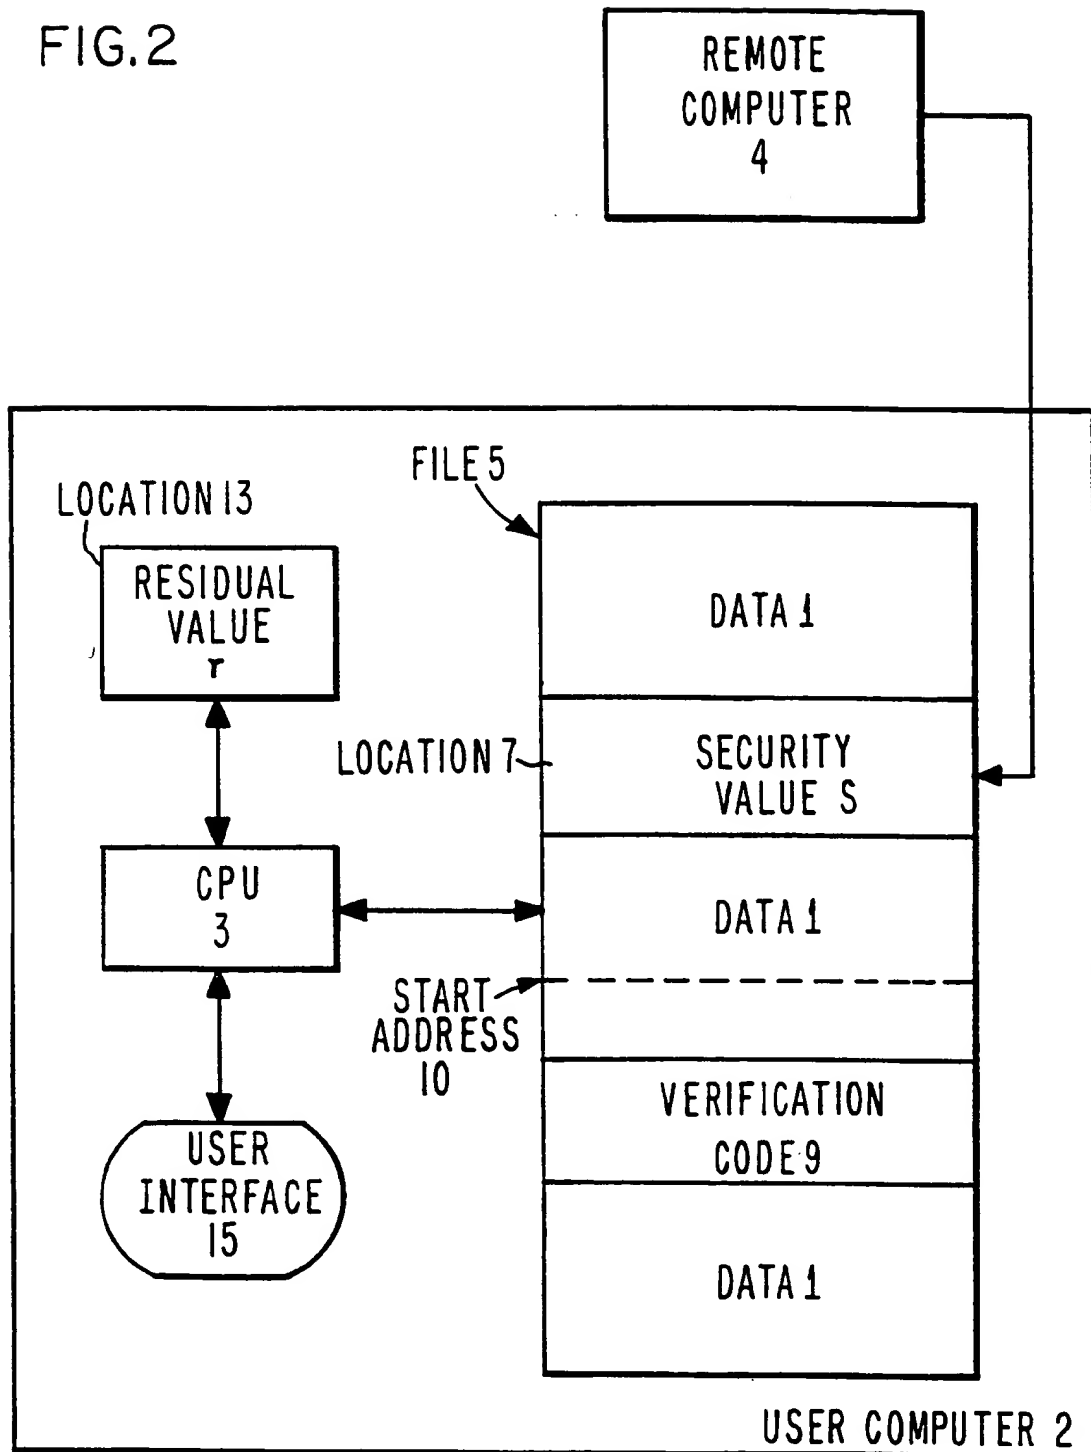
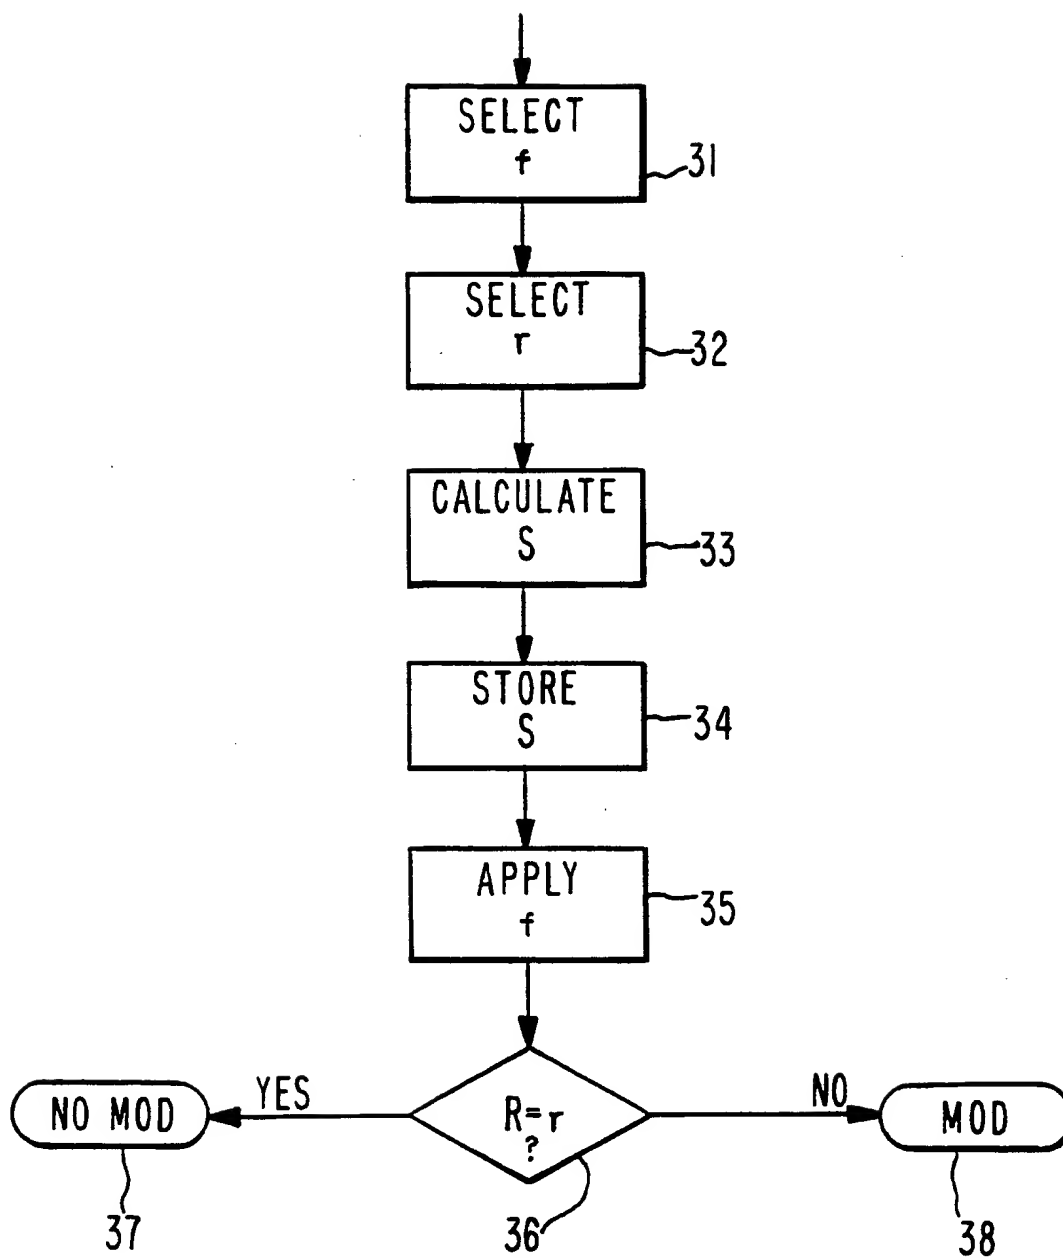


FIG. 3

3/3



INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 99/01444

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 G06F11/00 G06F1/00

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 93 25024 A (CYBERLOCK DATA INTELLIGENCE, INC.) 9 December 1993 see the whole document ---	1-11
A	WO 95 15522 A (SCHEELE ET AL.) 8 June 1995 see figure 4C -----	1-11

☐ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"Z" document member of the same patent family

Date of the actual completion of the international search

28 May 1999

Date of mailing of the international search report

07/06/1999

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Absalom, R

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 99/01444

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 9325024 A	09-12-1993	NONE	
WO 9515522 A	08-06-1995	AU 1182195 A	19-06-1995